102(b) as anticipated by U.S. patent 6,073,168, Mighdoll, et al., *Method for reducing delivery latency of an image or other secondary information associated with a file,* issued 6/6/2000, henceforth "Mighdoll". Claims 8, 13, 17, and 23 were rejected under 35 U.S.C. 103(a) as obvious in light of Mighdoll combined with U.S. patent 5,751,581, Tau, et al., *Material movement server,* issued 5/12/98, henceforth "Tau". Claims 9, 14, 18, and 24, finally, were rejected under 35 U.S.C. 103(a) as obvious in light of Mighdoll combined with U.S. patent application publication 2002/0091853 A1, Moore, et al. *Enhancing application performance in dynamic networks,* having a priority date of 1/5/2001, henceforth "Moore". Applicants are traversing the rejections.

Traversal

In the following, Applicants will first describe what Applicants are claiming and will then demonstrate that Mighdoll does not disclose all of the limitations of Applicants' claims 5-7, 10-12, 15-16 and 19-22, as required for a rejection under 35 U.S.C. 102, that the combination of Mighdoll with Tau does not disclose all of the limitations of Applicants' claims 8, 13, 17, and 23 as required for a rejection under 35 U.S.C. 103, and that the combination of Mighdoll with Moore does not disclose all of the limitations of claims 9, 14, 18, and 24, also as required for a rejection under 35 U.S.C. 103.

What Applicants are claiming

Applicants' claim 5 broadly sets forth Applicants' invention as a "method employed in a distributed database system . . . for responding to a request":

> 5. A method employed in a distributed database system that includes a plurality of database systems for responding to a request received in a database system of the plurality,
> the method comprising the steps performed during execution of the request in the database system of:
> determining that the request is preferably executed at least in part in another database system of the plurality; and
> redirecting the execution of at least the part of the request to the other database system.

Beginning with definitions of important terms, as may be seen from Applicants' FIG. 1, the "distributed database system" includes "a plurality of database systems". It is clear from the

discussion beginning at page 3, line 1 that a "database system" is a system that can respond to requests that include queries written in a query language such as SQL. It is further important to claim 5 that the steps of the method are performed "during execution of the request in a database system of the plurality". Finally, the method steps involve "another database system of the plurality", that is, another database system belonging to the distributed database system.

What Mighdoll discloses

As shown in FIG. 1, Mighdoll discloses a system for providing Web access via a WebTV server 5 to a number of WebTV clients 1. WebTV clients are systems in which the display device is an ordinary television set (col. 3, line 45-col. 4, line 48). The disclosure of Mighdoll that is of interest in the present context concerns WebTV server 5 and is found in FIG. 4 and FIG. 9.

FIG. 4 shows details of WebTV server 5. The description of the figure begins at col. 5, line 7. As set forth there, WebTV server 5 is a "caching proxy server", that is, it has a proxy cache 65 which contains copies of documents that have been accessed in the past by server 5's WebTV clients 1. WebTV clients 1 access the Web via WebTV server 5. If a copy of the document requested by a WebTV client 1 is present in WebTV server 5, server 5 provides the copy to the client; otherwise, WebTV server 5 fetches the document from the server on the Web in which the document resides and provides the document to WebTV client 1. Generally, WebTV server 5 will also place a copy of the fetched document in proxy cache 65.

A special feature of WebTV server 5 is shown in the flowchart of FIG. 9 and explained beginning at col. 12, line 59. As set forth there, the location in the Web specified by a URL for a document may not be the document's actual location; instead, the specified location may contain another URL which is a forwarding address to another location in the Web.. This forwarding address is termed a "redirect address" in Mighdoll. Usually, when what is at the location specified by a URL is a redirect address, the server indicated by the URL returns the redirect address to the requesting Web server (here, WebTV server 5) and that Web server then uses the redirect address to obtain the document specified by the original URL from the location specified by the redirect address. WebTV server 5 speeds this process up by caching redirect addresses in document database 61 as well as documents. As shown in FIG. 9, when server 5 receives a URL from a client (901) and there is a redirect address for the URL in database 61, server 5 uses the redirect address to retrieve the document (902,906); otherwise, server 5 uses the URL. If a

redirect address is received in response to either URL, the redirect address is saved in database 61 (904, 905).

Application of Mighdoll to claim 5

Applicants' claim 5 is directed to a

> A method employed in a distributed database system that includes a plurality of database systems for responding to a request received in a database system of the plurality (preamble)

As pointed out in the discussion of claim 5 *supra*, the claim term "database system" is defined in the Specification as a system that can respond to requests that include queries written in a query language such as SQL. A standard Web server responds to URLs, not to "requests that include queries written in a query language", and there is absolutely no indication in Mighdoll that either WebTV server 5 or the remote Web servers 4 shown in FIG. 1 are different from standard Web servers in this respect. To be sure, WebTV server 5 contains a document database 61 and a user database 62, but there is also no indication that either of these is a "database system" as that term is used in claim 5. The document copies, the detailed information about the copies (col. 6, lines 20-col. 7, line 30), and the redirection addresses in document database 61 are all organized by URL and the user information in user database 62 is most likely organized by some kind of user ID. Using a URL to fetch a copy of a document, detailed information about the copy, or a redirect address for the URL cannot reasonably be taken to be an "execution of a request" as the term "request" is used in claim 5, and nothing beyond using a URL to fetch information is required in the system of Mighdoll. There is thus no disclosure whatever in Mighdoll of anything that corresponds to Applicants' distributed database system.

Since there is no "distributed database system" disclosed in Mighdoll and no "database system of the plurality" that responds to a "request", and fetching a document specified by a URL cannot reasonably be termed an "execution of [a] request", the claimed method cannot be performed in the system of Mighdoll and the reference does not support the rejection under 35 U.S.C. 102(b). This analysis is confirmed when Examiner's rejection of claims 5-7, 10-12, 15, 16, and 19-22 under 35 U.S.C. 102(b) is examined in detail. Examiner finds the step of

> determining that the request is preferably executed at least in part in another database system of the plurality

as occurring when WebTV server 5 determines that there is a redirect address for the desired URL and the step of

> redirecting the execution of at least the part of the request to the other database system

as occurring when WebTV server 5 uses the redirect address to access the document on another server.

The rejection of course requires that WebTV server 5 and the other web servers 4 be taken as "database systems", which they are not, that the URL be taken as a request which involves a query in a query language, which it is not, and that responding to a URL be taken as an execution of a request in a database system, which it is not, but even if WebTV server 5 is taken to be a database system and a URL is taken to be a request, WebTV server 5 doesn't perform either claimed step. When a document is available in WebTV server 5's cache, server 5 responds to reception of the document's URL from a WebTV client by retrieving it from the cache, whether or not there is also a redirect address associated with the document. When a document is not available in the cache, server 5 responds by using the redirect address if it is present and otherwise the document's URL to retrieve the document from a remote server 4. The presence of a redirect address for a URL thus has nothing whatever to do with whether the document is retrieved from the cache or from a remote server, and step 902 of FIG. 9 does not represent a disclosure of the method step of

> determining that the request is preferably executed at least in part in another database system of the plurality.

Further, server 5 does not respond to the absence of a document in the cache by

> redirecting the execution of at least the part of the request to the other database system

as required by the claim, but instead simply sends either the URL originally received from the client or the redirect address, if there is one, to the Web server specified by the URL or the redirect address.

Since Mighdoll discloses neither the system in which Applicants' method of claim 5 is practiced nor the steps of the method, Examiner's rejection of claim 5 as anticipated by Mighdoll is without foundation. As Examiner will immediately see, the arguments that apply to claim 5

apply also to claim 15, and because those independent claims are patentable over Mighdoll, so are claims 6-14 dependent from claim 5 and claims 16-29 dependent from claim 15.

The rejection of claims 6, 11, 19, and 21

Continuing with the rejection of dependent claims 6, 11, 19, and 21 as anticipated by Mighdoll, claim 6 is exemplary for these claims. It includes the further limitation that

> the step of determining whether the request is preferably executed in the other database system determines that an object required for execution of the request is lacking in the database system.

In his rejection, Examiner refers Applicants to the discussion of the cached redirect address beginning at col. 12, line 59. As already pointed out, the cached redirect address has nothing whatever to do with "determining whether an object required fro execution of the request is lacking in the database system", as required to anticipate these claims. More to the point here are the facts that a URL cannot reasonably be termed a "request" as that term is defined in Applicants' Specification, that using a URL to fetch a document cannot reasonably be termed an "execution of a request", and that responding to a miss after an application of the URL to a cache by using the URL to fetch the document from a remote server cannot reasonably be termed a "redirection of the execution ... of the request...", all as required by the claim, and Mighdoll consequently does not anticipate claim 6. The same argument can of course be applied with regard to claims 11, 19, and 21.

The rejection of claims 7, 12, 16, and 22

The added limitation in claim 17 is

> placing the request in a form required for execution in the database system;
>
> modifying the form when it has been determined that the request is preferably executed at least in part in another database system; and
>
> in the step of redirecting, the modified form is redirected.

Examiner again refers us to the discussion of redirect addresses at col. 13, lines 3-14, but as pointed out above, in WebTV server 5, the presence of a redirect address for a URL has nothing whatever to do with whether the document is retrieved from the cache or from a remote server, and the cited location therefore does not disclose the limitations. Further, as would be expected by the fact that a URL cannot reasonably be termed a "request" as that term is used in

Applicants' claims, and the fact that fetching a document corresponding to the URL cannot reasonably be termed an "execution" of the URL, there is simply no disclosure at all in Mighdoll of anything corresponding to the steps of the added limitation, and thus Mighdoll does not anticipate claim 7. Again, the argument made with regard to claim 7 can be applied to claims 12, 16, and 22.

The rejections under 35 U.S.C. 103

The rejection of claims 8, 13, 17, and 23

Claim 8 is exemplary. The added limitation is this:

> the request includes an SQL statement;
> the form required for execution is a cursor; and
> in the step of modifying the form, the cursor is marked for redirection.

The additional reference, Tau, which Examiner has employed to make the rejection merely discloses that cursors can be declared and used in SQL. There is nothing whatever in Tau that discloses or implies that a cursor can be marked for redirection and then redirected to another database system for execution, as required by the claim when read in the context of the claims it is dependent from. Since that is so, Examiner has not made his *prima facie* case of obviousness and his rejection of claim 8 and of claims 13, 17, and 23 are without foundation.

The rejection of claims 9, 14, 18, and 24

Claim 9 is exemplary. The added limitation is this:

> the request includes a call to a procedure object; and
> in the step of modifying the form, the call is rewritten in the form required for execution as a remote procedure call directed to the other database system.

The additional reference, Moore, which Examiner has employed to make the rejection merely discloses what remote procedure calls are and how they work normally; there is no suggestion in the reference that a call to a procedure object in a request made in a database system may, when a portion of the request is redirected, be rewritten as a remote procedure call directed to the other database system, which is what the claim sets forth when it is read in the context of the claims it is dependent from. Thus, in combining Moore with Mighdoll, Examiner has not made his *prima facie* case of obviousness and his rejection of claims 9, 14, 18, and 24 are without foundation.